



Office de la propriété
Intellectuelle
du Canada

Un organisme
d'Industrie Canada

Canadian
Intellectual Property
Office

An Agency of
Industry Canada

*Bureau canadien
des brevets*
Certification

La présente atteste que les documents
ci-joints, dont la liste figure ci-dessous,
sont des copies authentiques des docu-
ments déposés au Bureau des brevets.

*Canadian Patent
Office*
Certification

This is to certify that the documents
attached hereto and identified below are
true copies of the documents on file in
the Patent Office.

Specification and Drawings, as originally filed, with Application for Patent Serial No:
2,418,729, on February 11, 2003, by **CATENA NETWORKS CANADA INC.**, assignee of
Tim Jenkins and Sejal Patel, for "Demand Accelerated IGMP Leave".

Agence certificatrice/Certifying Officer

February 10, 2004

Date

Canada

(CIPO 68)
04-09-02

OPIC CIPO

DEMAND ACCELERATED IGMP LEAVE

Background of Invention:

The text is written for multicast media distribution over ATM but the concepts apply to other technologies where there is an explicit association between the availability of a resource (a virtual circuit or bandwidth) and the ability to provide a requested service (IGMP group or digital video stream).

Figure 1 shows a typical digital video distribution system for delivering digital video over IP over ATM over DSL. Figure 2 shows a possible distribution system for delivering digital video over ATM over DSL. In both figures, a multicast capable Broadband Loop Carrier performs the multicast by connecting source media streams to the media VCs that are connected to the DSL modems. It is through the DSL modems that the STB or STBs request the video channels.

In the former diagram, there are potentially multiple independent Set Top Boxes (STBs) that extract the desired video stream from a LAN that connects them to the DSL modem.

In the latter diagram, the STB function is integrated with the DSL modem, and this box might also provide a LAN interface for network access.

In general, the number of TVs serviced on the customer premises is the same as the number of desired video feeds for which the customer is subscribed.

1/ Digital video can be delivered to customers using DSL lines. Typically, customers are provisioned with separate ATM virtual circuits (VCs) to carry the video stream, where one VC is usually provisioned per video stream subscribed to by the customer. Additional VCs are normally necessary for media stream control and other use.

2/ IGMP is often used to request specific video streams from the network. To do this, a Set-Top Box (STB) reports group membership, where the group is a specific video channel or stream.

3/ There are three primary methods of disconnecting a specific video stream from the VC that is carrying it between the Multicast Capable BLC and the DSL modem:

1. One is the normal “leave” as described by the IGMP specifications. In this case, video streams are not immediately disconnected from the VC carrying them when a STB sends an IGMP Leave message. In IGMPv2 and IGMPv3, this time is the expiration of timers started by the reception of the IGMP Leave message. In IGMPv1, there is no Leave message, so disconnection took place as described by third mechanism.
2. A second is sometimes known as “fast leave”. In fast leave, the video stream is immediately disconnected when the IGMP Leave message is received. This can be done in a proprietary manner or by setting the timers associated with reception of the IGMP Leave message to very short intervals with no retries.

3. The third method is a time-out based on the lack of appearances of Membership Reports for a particular group. This time-out interval is based on the number of times a General Query has been sent on the interface used by the video stream. This is the only method of disconnection when using IGMPv1.

Limitations and issues:

The IGMP protocol is designed to allow servers to be unaware of the exact number of clients that are members of a group. It is also designed such that group members do not report their own membership if they detect a peer in the same group. The result of this is that if more than one STB is receiving the same group (channel if video multicast), not all of the available media VCs are in use. In contrast, when all clients are using different groups, all available media VCs are in use.

Potential problems in the delivery of the video (or other multicast stream) can arise in either of these cases, depending on the “leave” mechanism used on the interface.

If the interface is assigned to fast leave and there is only a single media VC in use by multiple STBs, some subscribers can see a glitch in group reception if one of the STBs sends a leave. This is because the server stops delivering the group immediately. Service is not restored for the original group until those STBs that did not change groups report their membership.

If the interface is assigned to use normal leave behaviour and all media VCs are in use, a channel (or group) change request can go unanswered due to the lack of availability of media VCs. In this case, a typical user changes channels. This is done by the STB sending an IGMP Leave message followed by an IGMP Report for the new channel. However, due to normal IGMP leave behaviour, the group just left has not been disconnected from the media VC. Since all other media VCs are in use, the report for the new group is dropped.

When other methods of group (channel) delivery are used, different resource limitation issues can arise. For example, if a single ATM VC is used but provisioned with sufficient bandwidth for the appropriate number of media stream subscribed, there can be periods of time for which more bandwidth is required than the VC has been provisioned for. The bandwidth issue is equally applicable to all media paths going to the STB or STBs. At this time, all standards bodies have rejected this particular implementation of transporting video over ATM. Regardless, the concepts here are applicable to any resource usage by groups in general, and not to just video delivery.

A number of possible solutions to this problem have been proposed. However, the method described herein is typically easier to administer and has good results under most conditions.

Other proposals to solve this problem include:

- 1) Always provision one more media VC than there are STBs and use slow leave. This method increases the administrative requirements and requires the permanent use of a media VC for transient conditions. For delivery mechanisms other than one ATM VC per subscribed group, these would be the same as provisioning extra bandwidth for all customers. Phrased another way, this means $N+1$ times the bandwidth is needed for N TVs.
- 2) Keep track of individual members in a group. This method is not scalable, and does not solve the problem if some members do not report their membership in a group because they detect peers already in the same group.

Details of the Invention

This document describes a “demand accelerated leave” mechanism that solves the problems of video access in both cases described above.

In “demand accelerated leave”, the reception of an IGMP Leave message triggers the normal (slow) leave behaviour. In this case, a Group Specific Query is transmitted back to the clients a certain number of times at a certain interval. The purpose of this is to solicit responses from members of the group in order to determine that the group should not yet be disconnected from the media VC. If no IGMP report is received for that group by the end of the retries, the group is disconnected from the media VC. If an IGMP report is received for that group, the timers are cleared and the group is not disconnected.

When IGMP reports are received for a group that isn’t already being sent to a port, the server must look for an available media VC (or other resource) on which to send the group. As described above, there are conditions under which no media VC (or other resource) is available.

In “demand accelerated leave”, the counter and timer associated with the Group Specific Queries are used to determine if a group already being sent to a port is aging. Aging means that a Group Specific Query has been sent and there exists the possibility that the group might be disconnected from the port in the near future. Further, it is proposed that the probability that a group will be disconnected is proportional to the time between the transmission of the first Group Specific Query and the current time. This means that the group with the oldest Group Specific Query timers is the most likely to be disconnected.

“Demand accelerated leave” then causes the oldest group to be disconnected from its media VC in order to provide an available media VC for the group requested in the just received IGMP report. In effect, it turns a “slow” leave into a “fast” leave based on a reasonable heuristic when the demand for a media VC requires it.

The following pseudo code illustrates the use of this mechanism when an IGMP membership report arrives:

```
for ( each media VC assigned to the interface )
{
```

```

if ( an available media VC has not been found AND
    this media VC is available )
{
    save this media VC as the available media VC
}
else if ( this media VC is not available )
{
    if ( this media VC is carrying the requested group )
    {
        cancel Group Specific Query timers
        reset General Query counter
        done
    }
    else if ( this media VC's group is aging AND
              it's older than the oldest )
    {
        set this media VC as the oldest
    }
}
}

if ( there is an available media VC )
{
    use it
}
else if ( there is a media VC marked as oldest )
{
    stop using this media VC for the group that's currently on it
    use it
}

```

The algorithm to determine if a media VC is oldest can be based on the age of the Group Specific Queries, or can be further enhanced by looking at the number of missing responses to General Queries if no media VCs are marked as old based on the Group Specific Query timers.

While this algorithm as described is specific to multicast transmission over ATM controlled by IGMP, the concepts could be made more generic to describe how to select any limited resource based on service requests. The generic algorithm is:

```

request for service arrives
if ( insufficient resources are available for service )
{
    for ( each service currently being provided )
    {
        find oldest/most likely to be unneeded
    }
}

if ( old/unneeded service found )
{
    stop that service
    start requested service
}

```

Using this generic mechanism, the algorithm described above might be able to be extended to other resources than ATM VCs and other services than multicast groups. Note that resource limitations can also consist of bandwidth availability limits in addition to the number of VCs available.

However, other control protocols than IGMP may have completely different algorithms to determine which of any service currently being provided is the oldest.

This mechanism also limits the bandwidth usage of a group delivery system to that of the bandwidth requirements of the number of provisioned groups, even in the presence of a non-bandwidth limited media. This makes bandwidth usage more efficient.

Claims:

1. A mechanism of determining the oldest of groups that have a probability of being disconnected from a port when using IGMP as the control mechanism.
2. A mechanism for providing glitch-free multicast delivery for multiple group receivers when the number of available media streams (ATM virtual circuits or other media streams) is limited to the maximum number of groups that can be received on a port.
3. Implementation of 2) above by the use of 1) above.
4. A mechanism for limiting the bandwidth requirements of multicast group delivery to the number of requested groups.
5. An implementation of 4) using 1) above.
6. A mechanism for providing glitch free multicast delivery for digital video transmitted over ATM when there is one video channel delivered per VC.
7. An implementation of 6) using 1) above.
8. A mechanism for providing glitch free multicast delivery for digital video transmitted over ATM when there is one VC provisioned for the exact number of provisioned channels.
9. An implementation of 8) using 1) above.
10. A mechanism for providing glitch free multicast delivery for digital video (or any other service) transmitted over a bandwidth limited channel when the channel bandwidth is limited to the provisioned service capability.
11. An implementation of 10) using 1) above.
12. A mechanism for limiting the bandwidth requirements of digital video delivery to the bandwidth requirements of the number of provisioned video channels.
13. An implementation of 12) using 1) above.
14. A mechanism for limiting the number of media streams required for glitch free delivery of digital media.
15. An implementation of 14) using 1) above.

Figure 1: Distributed STBs

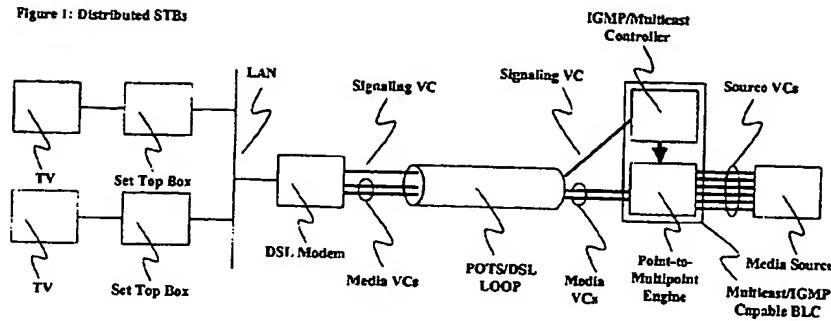
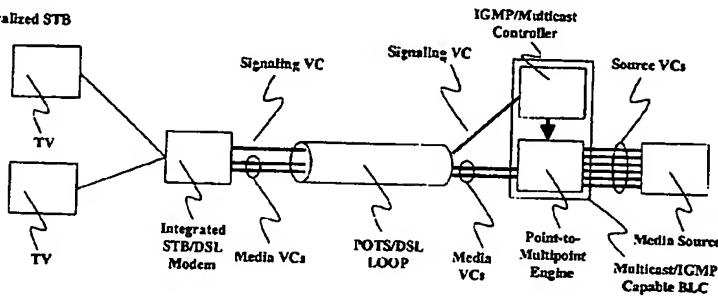


Figure 2: Centralized STB



Priority Document

Ardeshir Tabibi, Reg. No. 48,750 (650) 326-2400
Application No.: 10/773,085 Filed February 4, 2004
Method For Selecting A Resource To Provide A Requested Service In A
Multicasting Environment
Applicants: Jenkins et al.